Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

An Empirical Analysis of Digital Music Bundling Strategies

Brett Danaher, Yan Huang, Michael D. Smith, Rahul Telang

MANAGEMENT SCIENCE
Vol. 60, No. 6, June 2014, pp. 1413–1433
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.2014.1958

An Empirical Analysis of Digital Music Bundling Strategies

Brett Danaher
Department of Economics, Wellesley College, Wellesley, Massachusetts 02482, bdanaher@wellesley.edu

Yan Huang
Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109, yphuang@umich.edu

Michael D. Smith, Rahul Telang
School of Information Systems and Management, Heinz College, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213 {mds@andrew.cmu.edu, rtelang@andrew.cmu.edu}

We use panel data on digital song and album sales coupled with a quasi-random price experiment to determine own- and cross-price elasticities for songs and albums. We then develop a structural model of consumer demand to estimate welfare under various policy relevant counterfactual scenarios. This approach represents an early application of the "big data" management paradigm within the media industries and provides managers with detailed guidance on optimal pricing and marketing strategies for digital music. Our results show that tiered pricing coupled with reduced album pricing increases revenue to the labels by 18% relative to uniform pricing policies traditionally preferred by digital marketplaces while also increasing consumer surplus by 23%. Thus, optimal tiered pricing can yield a Pareto improvement over the prior status quo. Additionally, our results indicate that even without tiered pricing, unbundling albums outperforms "album-only" pricing policies that dominated the era of physical CD/cassette sales.

Keywords: information goods; bundling; digital music; price elasticity; natural experiment
History: Received September 18, 2012; accepted January 12, 2014, by Alok Gupta, special issue on business analytics. Published online in *Articles in Advance* May 5, 2014.

1. Introduction

Digital channels create many new challenges and opportunities within the media industries. These include changes to the media products themselves creating new interactions across products and channels (Zhu and MacQuarrie 2003), increased opportunities to use microlevel data to directly measure consumer behavior, new advances in statistical models to generate business intelligence from these data, and new opportunities to make direct changes to marketing strategies in response to this business intelligence. The goal of this paper is to illustrate how business analytics can help media firms address these challenges and opportunities by demonstrating how newly available consumer data, online marketing experiments, and advanced analytical tools can be used in the context of pricing and mixed bundling in the music industry.

Specifically, in our research setting, we worked with a major music label to develop a pricing experiment and associated empirical models to better understand how the label should set prices given the interrelated demand structure between albums and the component songs on those albums. In our experiment we systematically increased the digital prices of 2,000 best-selling singles by $0.30. The digital setting allowed us to directly set prices, to control the exact date and time the price change occurred, and to reliably measure the resulting sales of the singles under observation, the sales of all the other singles in the album, and sales of the album itself.

The digitization of music and sales channels was instrumental to producing the insights we contribute with this study for three reasons. First, the lower menu costs associated with digital sales channels greatly reduce the cost associated with price experimentation—the price of products can be quickly changed with a few mouse clicks rather than relabeling hundreds or thousands of CDs in inventory in a brick-and-mortar setting. Economic theory predicts that retailers will only change prices when the benefit from the price change outweighs the menu costs associated with making the price change, something that is much more likely online than in a store. Second, digitization has resulted in a partial disintermediation of the retailer—generally when a music label changes its wholesale price for a song or album to an online store the price change is immediately reflected in the retail price to the consumer.[1] This is in direct contrast to the brick-and-mortar world where wholesale price changes may have resulted only in very delayed changes to the retail price (if at all), making it hard for labels to measure the effect of price on consumer demand. Finally, digitization has enabled the more accurate and rapid reporting of sales data, allowing the wholesaler to observe changes in demand in real time and potentially attribute them to the precise timing of a price change, thus increasing potential returns to investing in price experimentation.

When one looks at these operational challenges, it is not surprising that prior to the development of digital channels, the major label we worked with on this study had not conducted any pricing experiments of this sort in any of its brick-and-mortar channels: The operational challenges were too high, and the resulting data were too noisy, and thus the value of the business intelligence to be gained did not justify the costs. Digitization has changed these factors, allowing media firms to easily conduct pricing experiments, reliably collect data on how consumers respond, use advanced modeling techniques to analyze these data, and take direct action based on the resulting business intelligence. Our goal in this paper is to highlight "big data" as a management paradigm transforming even the media industries, where microlevel data sets may be relatively small, but where digitization has rendered online experiments and advanced data analytics techniques into powerful and viable tools to improve managerial decision making through data and evidence.

Our specific empirical question explores how music labels should optimize pricing strategies in the presence of bundled (album) and unbundled (single) content. A key empirical challenge we face in this setting is that a price change to an individual song will affect not just sales of that song, but also sales of both the album and of other songs. A change in the price of a song affects demand for its parent album through changing the utility of purchasing the album (bundle of songs) relative to just purchasing the individual song. Moreover, this change in album demand affects demand for other songs on the album indirectly. Because of this, one cannot cleanly identify the overall effect of changes to album and each single song prices from reduced-form models, especially in the cases where the prices of several singles (or the album) increase, because it is hard to tease apart the net effect of each price change. Instead, we need a model that captures these dynamics in a parsimonious manner. We provide this in our study by constructing a structural model of demand, which allows us to capture consumers' choices between individual digital single songs and albums (which is a bundle of the individual songs) and allows us to simulate the welfare outcomes of various counterfactual pricing policies. We then estimate our model with data from a plausibly exogenous pricing experiment[2] to understand consumers' price sensitivity and the mean value distribution of songs within a typical album.

Using these parameter estimates, we numerically solve the pricing problem for an "average album," as well as for albums of different genres and different popularities. Our results show that demand for the most popular songs is relatively inelastic (suggesting higher optimal prices for popular tracks), demand for albums is elastic at the preexperimental album prices (suggesting lower optimal prices for albums), and tracks and albums are weakly substitutable for each other. Our simulation results also show that record labels, artists, and copyright holders can receive significantly more revenue under tiered pricing strategies in digital marketplaces versus the uniform ($0.99) pricing policies historically offered by many marketplaces. Importantly, consumer surplus also increases significantly under tiered pricing strategies when albums are adjusted to optimal prices—and thus our results suggest a Pareto-improving scenario whereby both producers and consumers gain from increased price flexibility along with optimization. Finally, our simulations show that, even under the traditional uniform pricing regimes, unbundled sales not only increase consumer surplus, but also provide more revenue to the artists and labels than would be available under an "album-only" sales policy.

2. Literature Review

The market for recorded music is one of the biggest segments of the entertainment markets. Prior to digitization, there had been few studies investigating the characteristics of demand for music, probably because of the difficulty in obtaining detailed price and sales data discussed above. For example, one notable paper on the industry (Belinfante and Johnson 1982) theorizes that demand for music is relatively inelastic, and as such competition centers around product differentiation

[1] Importantly, this is not a legal requirement. Online retailers do not have to change their retail prices in response to wholesale price changes, nor have most online retailers contractually obligated themselves to do so. Nevertheless, the status quo in the industry is that major online retailers have historically allowed the retail price to automatically change with the wholesale price.

[2] As we will describe in detail, the pricing experiment was not a fully randomized trial, but price changes were implemented according to a rule, leaving us with a potential source of some exogenous variation in prices. We will be very clear about both the strengths and potential weaknesses of our pricing experiment, but we note that a major contribution of our paper is in working with pricing/sales data from a major music label involving their top 2,000 songs rather than a small, unrepresentative group of unpopular songs. Importantly, it is very unlikely that a major music label would every truly randomize pricing for their 2,000 most popular tracks, and our pseudoexperiment represents the next best alternative.

rather than price. However, likely owing to a lack of data, we are aware of no papers from 1982 until recently that empirically test this hypothesis.

Moreover, even these recent studies have primarily relied on consumer surveys as opposed to market data. For example, Breidert and Hahsler (2007) introduce three approaches based on adaptive conjoint analysis to estimate consumer's willingness to pay for bundles of music downloads, finding that marginal willingness to pay per title decreases with larger package sizes. Similarly, Shiller and Waldfogel (2009) explore the firm's revenue under various pricing schemes based on their surveyed willingness to pay results. Their study suggests that alternative pricing schemes, including simple schemes such as pure bundling and two-part tariffs, are revenue improving.

In contrast, Elberse (2010) represents one notable study in the literature to use market data. Elberse (2010) applies system-of-equations modeling methods to examine the effect of digital unbundling on album sales (both physical and digital) and finds that "revenues decrease significantly as digital downloading becomes more prevalent, but the number of items included in a bundle (a measure of its 'objective' value) is not a significant moderator of this effect" (p. 107). However, this study does not account for the effect of the prices of the singles and albums because there is no price variation in the data set. Moreover, it compares the unbundled *digital* marketplace to the bundled *physical* marketplace, complicating the interpretation of the resulting estimates. In contrast, one finding of our paper compares the unbundled digital marketplace to a counterfactual pure bundled digital marketplace (a comparison that may be more relevant to decision making in the music industry).

Our paper also extends the current literature on digital music bundling by using a rich real-world data set combined with a quasi-randomized pricing experiment to empirically estimate the value distribution of songs in an album and the value of the album as a whole. Based on our structural estimates, we simulate optimal pricing strategies that maximize the joint profits made from single song sales and album sales.

In terms of theoretical foundation, our research relates most closely to the literature on bundled pricing strategies, and specifically to the literature discussing bundling problems in the context of information goods. The academic literature has long understood that appropriate bundling strategies can help monopolists increase revenue (e.g., Burstein 1960, Stigler 1963, Adams and Yellen 1976, McAfee et al. 1989, Armstrong 1996). Extending this literature, there are a variety of analytic models analyzing pricing and bundling strategies in the context of information goods—goods with high fixed costs to provide but low marginal costs to reproduce (e.g., Bakos and Brynjolfsson 1999, Geng et al. 2005, Hitt and Chen 2005). Most of these models show that some form of bundling is preferable to pure unbundling for information goods.

However, there are only a few empirical papers we are aware of that deal with pricing of product bundles using real-world price and sales data, and none in the context of digital music. In the context of pricing academic journals, Chuang and Sirbu (1999) discuss the bundling problem journal publishers face and employ both analytical and empirical methods. They find that in a context where consumers value only a subset of the bundle components, mixed bundling strategies dominates other strategies and component pricing outperforms pure bundling. Chu et al. (2006) estimate a structural demand model for eight plays showing at a Palo Alto theater. They find that all alternative bundling strategies yield higher revenue than a uniform pricing strategy does. Crawford (2008) estimates demand for TV bundles, finding that whereas bundling general interest cable networks has no discriminatory effect, bundling an average top special-interest cable network significantly increases the estimated elasticity of cable demand and raises profits. Crawford and Yurukoglu (2012) use a structural model approach to study the effect of bundling television channels on short-run welfare. They use historical data to calibrate their model and find that if input costs are fixed at current levels, unbundling TV channels can increase consumer surplus but reduce producer surplus. However, in practice, input costs will increase significantly under à la carte pricing, and the increase in input cost can offset consumer benefits from purchasing individual channels. Derdenger and Kumar (2013) examine bundling of consoles and software in the video game industry by incorporating the dynamic nature of durable good purchases. They find that pure bundling performs significantly worse than mixed bundling. Unlike these papers, we not only directly estimate the value distribution of components, but we also combine structural modeling with experimental data. This allows us to make more robust causal claims on estimated elasticity.

We use market-level demand data for album and song sales. Thus, in that regard, our method borrows from Berry et al. (1995; hereafter, BLP) and Berry (1994) and its probit extension proposed by Chintagunta (2001). The BLP estimating method has been widely adopted, notably by Nevo (2000, 2001) in the context of merger effects and market power, respectively, in the ready-to-eat cereal market, and by Davis (2006) in the context of movie theater markets. However, it is worth pointing out that our model is not a direct application of the BLP model. Although our model borrows a few high-level ideas from BLP about estimating individual choice model using market-level data, we are tackling a unique bundling setting, in which consumers choose individual songs versus albums, goods that are not necessarily differentiated (as they have same characteristics such as genre, artist, and so on). We will elaborate our model in §4.

3. Research Context

To empirically investigate the digital music bundling problem, we obtained a rich data set from a large anonymous record label, hereafter called "Label X." Label X sells its music through an online digital music retailer called "Retailer Y." In the data set, we observe weekly prices and sales at Retailer Y for many albums and their associated songs. The data set extends from January 2009 to July 2009 and includes a major price shift: Prior to April 2009, Retailer Y required that all songs sold on their digital marketplace be priced at $0.99. Starting on April 7, 2009, Retailer Y began offering three pricing tiers for songs: $0.69, $0.99, or $1.29 (Smith 2009). Because all of the data were provided by one label, we had access to all information regarding promotional events, campaigns, and other marketing shifts around these songs during our time frame. In the data, we see no special marketing activity outside of normal procedures for these songs.

Label X, wanting to study the impact of this price shift, selected a set of their most popular songs to undergo a wholesale price change resulting in retail prices shifting from $0.99 to $1.29.[3] Specifically, the design of the experiment was such that from April 17 onward, the top 200 selling songs at Label X (as measured by first quarter total sales at Retailer Y) were increased in price to $1.29.[4] Two weeks later, the prices of the 201–400th most popular songs were raised, then the 401–600th, followed by the 601–1,000th most popular songs. Finally, the 1,001–2,000th ranked songs were increased in price. It is important to note that when the price of a single increased to $1.29, it remained at that level until the end of the study period. Clearly, an ideal pricing experiment would involve random selection of songs into the treatment and control groups, but it is unlikely that a major label would ever be willing to randomly shift prices on a large set of their most popular songs. A strength of our data set is that it comes from a major music label and includes their 2,000 most popular songs, lending our results a high external validity that would not be present in a smaller or less representative setting. Moreover, the selection process for the treatment group is rule based and not based on future projected sales or trends, and as such the selection of songs that underwent a price increase may be random with respect to time (something we will explicitly test in the next section, finding that demand for these groups of tracks trended similarly over time until the price changes).

We only include albums that had at least one song that experienced a price change. We remove all "holiday" albums to eliminate confounding effects from seasonality in our estimations, and we remove all albums with more than 20 songs because these albums tend to be either compilations or collections of many niche titles and do not add much to our identification. Finally, we eliminate albums that are more than 10 years old because these albums are almost exclusively successful "classic" albums, which may systematically differ from other albums in our sample. After eliminating these songs and albums, we are left with 286 albums comprising 3,709 songs where 1,457 of these songs went through the price change.

Tables 1 and 2 provide song- and album-level summary statistics for our sample, and Table 3 summarizes the counts of albums by genre. These tables show that the majority of the time songs are priced at $0.99, and that only a segment of songs were treated with the price change and only at some date after April 17. The variation in album price is larger, but more than half the albums in our sample were priced at $9.99.

Table 1 Summary of Song Attributes

	Mean	Std. dev.	Max	Median	Min
Song price (per week)	1.02	0.09	1.29	0.99	0.99
Song sales (per week)	148.52	1,538.84	126,100	12	0
Week since release	234.98	131.09	520	203	32

Table 2 Summary of Album Attributes

	Mean	Std. dev.	Max	Median	Min
Album price (per week)	10.33	1.85	16.99	9.99	0.99
Album sales (per week)	75.52	253.28	8,295	27	0
Week since release	235.15	129.71	520	204	32
No. of songs in album	12.91	3.13	19	13	3
Catalog (new release) (%)		88.5 (11.5)			
Best of (not best of) (%)		26.2 (73.8)			

Table 3 Summary of Album Attributes: Genre Distributions

Genre	Number of albums	Genre	Number of albums
Alternative	79	New age	1
Blues	1	Other (kids/video/etc.)	7
Christian	7	Pop/rock	120
Classical	8	Rap	12
Country	29	Urban	8
Jazz	3	Vocals/lounge	3
Latin	1	World	7

[3] Label X did not adopt the $0.69 price point for any of their songs during our sample period, but later experimented with this price.

[4] For this and all future descriptions, it is legally important to note again that Label X cannot choose the retail prices. Rather, Label X chose to shift some songs to the highest wholesale pricing tier, resulting in an increase in the price at Retailer Y to $1.29. For brevity, in the rest of this paper we will simply refer to price increases to $1.29.

Figure 1 Average Weekly Sales by Rank (Prior to the Price Experiment)

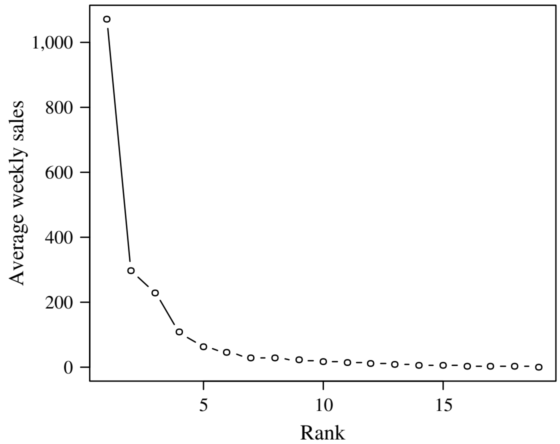

We rank songs in an album based on their sales prior to the price experiment (i.e., the highest-selling song prior to the experiment is ranked the first, the song that has the second-highest sales prior the experiment is ranked the second, and so on). Figure 1 shows the heterogeneity in song sales in the same album. This figure suggests that, on average, sales of the most popular song in each album are much larger than sales of other songs in the same album. Sales of songs in ranks 1–5 drop very quickly, and the difference in sales of songs that ranked lower than 5 is not significant. We discuss how we address these empirical characteristics in our estimation section.

3.1. Treatment vs. Control Group

In our sample, some songs went through a price change, whereas the others did not. We plot the sales trends of these two groups in Figure 2. We note that the trends of sales prior to week 17 are very similar, and that once the experiment starts in week 17, we see an immediate divergence in sales whereby songs that experienced a price increase have significantly lower relative sales than other songs do. This divergence increases over time as more songs in the treated group begin to undergo their price changes. Thus the untreated songs seem to be a good control for treated songs.

Instead of plotting absolute sales, we plot the average percentage change in sales from week to week in

Figure 2 Average Weekly Log Sales

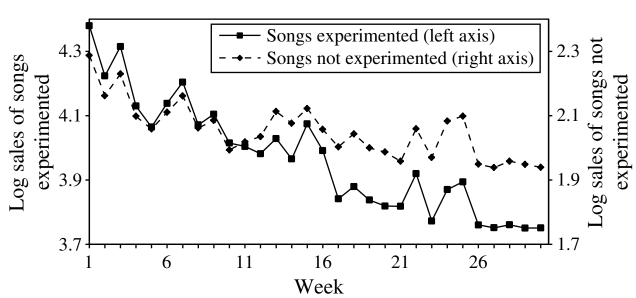

Figure 3 Average Percentage Change in Sales from Week to Week for (a) All Songs and (b) Experimented Songs

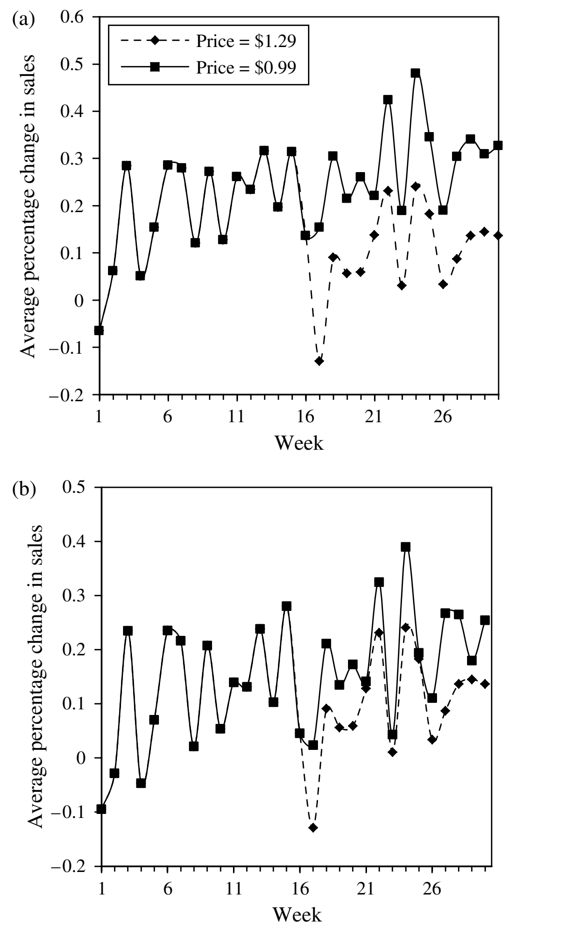

Figure 3. The "price = $0.99" group consists of all songs for which the price is $0.99 in the corresponding week (including songs whose prices never changed and the songs whose prices have not yet changed), whereas the "price = $1.29" group consists of songs for which the prices is $1.29 in the corresponding week. Note that the composition of $1.29 price group changes each week, as more songs join the group when their prices get changed. As we can see from Figure 3(a), sales of songs decline significantly as soon as the song price increases.

However, a confounding factor in Figure 3(a) is that the untreated songs had significantly lower sales to begin with, such that while the trends prior to experiment are similar, we may still worry about whether they are a good control. As a robustness check, recall that different songs went through price changes at different times. Thus, we can consider only those songs that went through the price change in part of our analysis and exploit time variation in the experimental window as a way to identify the effect of the price change. In this case the songs that went through the price change are more homogeneous and hence may

Table 4 Summary of Regression Results of β_{st}

	s												
t	17	18	19	20	21	22	23	24	26	27	28	29	30
2	−0.094	−0.057	−0.007	−0.104	−0.202	−0.022	−0.350	−0.047	−0.140	0.073	0.818	0.036	0.032
3	−0.088	−**0.368**	0.018	−0.049	−0.329	0.010	0.007	0.016	−0.039	0.113	**1.731**	0.111	0.008
4	−0.165	−0.106	−0.025	−0.156	−0.391	−0.002	−0.127	−0.037	−0.072	0.060	**1.694**	0.091	−0.059
5	−0.218	−0.547	−0.088	−0.160	−0.305	−0.050	−0.013	−0.084	−0.102	0.222	1.076	0.076	−0.016
6	−0.189	0.123	−0.006	−0.117	−0.134	−0.066	−0.150	−0.011	0.063	0.115	0.793	0.158	−0.036
7	−0.171	**0.280**	0.037	−0.101	−0.310	−0.011	−0.044	−0.052	0.038	0.352	0.423	0.135	−0.009
8	−0.147	0.106	−0.001	−0.153	−0.293	−0.074	0.153	−0.017	0.042	0.182	0.651	0.171	−0.017
9	−0.130	−**0.395**	−0.030	−0.127	−0.339	−0.062	−0.026	−0.039	−0.020	0.529	0.727	0.133	0.027
10	−0.199	−**0.858**	−0.042	−0.115	−0.447	−0.067	0.050	0.031	0.012	−0.069	0.463	0.015	0.052
11	−0.237	−0.213	−0.097	−0.155	−**0.512**	−0.072	−0.131	0.005	−0.019	0.212	0.985	−0.052	0.042
12	−0.247	−0.294	−0.101	−0.130	−**0.464**	−0.100	0.040	−0.012	0.105	0.303	0.901	−0.282	0.008
13	−**0.296**	−**0.361**	−0.163	−0.099	−0.254	−0.087	−0.247	−0.023	0.003	0.289	0.384	0.192	0.000
14	−**0.295**	−**0.343**	−0.157	−0.167	−0.223	−**0.160**	−0.340	−0.046	−0.164	0.066	0.669	**0.384**	−0.028
15	−0.206	−0.215	−0.145	−0.070	−0.326	−0.040	−0.216	0.017	−0.140	−0.283	1.067	0.002	0.047
16	−0.256	−0.280	−0.118	−0.090	−**0.483**	−0.062	−0.347	−0.002	−0.216	−0.007	0.443	−0.099	0.076
17	−**0.443**	−**0.364**	−**0.170**	−0.163	−0.417	−**0.143**	−0.322	−0.107	−0.203	0.111	0.495	−0.219	−0.037
18	−**0.474**	−**0.531**	−0.198	−0.151	−**0.470**	−**0.155**	−0.005	−0.093	−0.196	0.239	0.300	−0.180	−0.030
19	−**0.436**	−**0.546**	−**0.260**	−0.115	−0.329	−0.128	−0.284	−0.088	−0.246	0.291	0.635	−0.008	−0.069
20	−**0.406**	−**0.439**	−**0.306**	−**0.267**	−0.006	−0.114	−0.148	−**0.115**	−0.290	0.277	0.358	−0.209	−0.034
21	−**0.341**	−**0.449**	−**0.225**	−**0.226**	−0.399	−0.118	−0.120	−0.084	−0.175	0.354	0.792	0.083	−0.026
22	−**0.310**	−**0.394**	−**0.163**	−**0.168**	−0.230	−**0.263**	−0.115	−0.046	−0.172	0.347	**2.208**	0.249	−0.036
23	−**0.365**	−**0.408**	−**0.278**	−**0.235**	−0.347	−**0.226**	−**0.764**	−0.085	−0.174	0.250	0.381	−0.006	−0.058
24	−**0.346**	−**0.359**	−**0.255**	−**0.245**	−**0.469**	−**0.234**	0.084	−**0.216**	−0.134	0.506	−**1.523**	−0.052	−0.008
25	−**0.309**	−**0.417**	−**0.159**	−**0.161**	−**0.515**	−**0.221**	−0.220	−**0.208**	−0.141	0.391	−**1.536**	−**0.421**	0.000
26	−**0.307**	−**0.528**	−**0.142**	−**0.133**	−0.240	−**0.183**	−0.173	−**0.191**	−**0.425**	**0.694**	−1.388	−0.276	−0.021
27	−**0.314**	−**0.568**	−**0.193**	−**0.180**	−0.346	−**0.172**	−**0.480**	−**0.183**	−**0.566**	−0.074	−0.691	−0.222	−0.006
28	−**0.357**	−**0.600**	−**0.179**	−0.125	−0.417	−**0.159**	−0.237	−**0.222**	−**0.566**	−0.081	−**1.401**	−0.028	−0.020
29	−**0.299**	−**0.518**	−**0.183**	−0.123	−**0.453**	−**0.186**	−0.235	−**0.222**	−**0.508**	−0.022	−**1.396**	−**0.631**	−0.031
30	−**0.312**	−**0.540**	−**0.195**	−0.110	−0.363	−**0.169**	−**0.320**	−**0.199**	−**0.468**	0.701	−0.693	−**0.347**	−**0.115**

offer a more robust sample. Figure 3(b) shows that sales of this more homogeneous group of songs still decline sharply as soon as songs undergo the price change. This confirms our intuition that price change is the key driver for shifting demand. However, the difference between the two lines is slightly smaller.

To more formally test whether the control group is suitable for our analysis, we run the following regression:

$$LogSales_{ajt} = c_{aj} + \sum_{t=2}^{30} d_t D_t + \sum_{s=17}^{30} \sum_{t=2}^{30} \beta_{st} C_{ajs} D_t + \varepsilon_{ajt}.$$

Here, subscript a is the index for album, and subscript j is the index for rank. A combination of a and j uniquely identifies a song. The song-specific intercept c_{aj} captures song fixed effects, and ε_{ajt} captures any unobserved demand shock. Time dummy variable D_t represents a set of time dummies, which equals 1 in week t and 0 in other weeks; and its corresponding d_t is the week-specific effect for week t. Dummy variable C_{ajs} indicates the price of song j in album a was increased in week s. For example, if the price of this song was increased in week 20, then C_{ajs} equals 1 for $s = 20$, and 0 for any $s \neq 20$. Note that C_{ajs} is song specific but time invariant. For each combination of s and t, β_{st} is the coefficient of the interaction term $C_{ajs}D_t$. For the song we mentioned above, in week 2, $D_2 = 1$ and $D_t = 0$ for any $t \neq 2$. Therefore, $C_{ajs}D_t = 0$, as long as $t \neq 2$ or $s \neq 20$. In other words, the only nonzero $C_{ajs}D_t$ occurs when $s = 20$ and $t = 2$, and $LogSales_{aj2} = c_{aj} + d_2 + \beta_{20,2} + \varepsilon_{aj2}$. In week 20, $D_{20} = 1$ and $D_t = 0$ for any $t \neq 20$. Therefore, $C_{ajs}D_t = 0$, as long as $t \neq 20$ or $s \neq 20$. In other words, the only nonzero $C_{ajs}D_t$ occurs when $s = 20$ and $t = 20$, and $LogSales_{aj20} = c_{aj} + d_{20} + \beta_{20,20} + \varepsilon_{aj20}$.

If the songs that did not experience the price change are good controls of the songs that did experience the price change, we would expect β_{st} to be insignificant if $t < s$, which means that before the price change occurs, the sales of a song whose price will be changed later should not be statistically different from the sales of other songs.

Table 4 summarizes the estimates of β_{st} for each pair of s and t. Note that the column in which $s = 25$ is not included because none of the songs entered the experiment window in week 25. The estimates in bold are significant at the 0.05 level. As expected, with a few exceptions, most β_{st} with $t < s$ (above the gray diagonal line) are insignificant, whereas most β_{st} with $t \geq s$ are significant and negative. These results confirm our graphical intuition above that before the price change occurs, sales of songs that underwent the price change and sales of songs that did not undergo the price change are similar, suggesting that the experimental price change is causing any observed price divergence.

Moreover, as we will show later in our main analysis, the control group consists of not only songs that were not selected to go through the price experiment, but also songs whose prices would eventually increase but have not been treated in the current period. For example, in week 20, songs that entered the experiment in week 17 to week 20 form the treatment group, and songs that entered the experiment after week 20, as well as songs whose prices never increased, form the control group, and the difference between the two groups contributes to the identification of the effect of price change. Since β_{st}'s have similar trends with t across s (the week in which the price increase occurred), including "not yet treated" songs as controls can help control for additional heterogeneity in sales trends between songs that eventually went through the experiment and the ones that never experienced the price change. In addition, we can also identify the effect of the price change by exploring variations in the time when the price change occurred within singles that eventually went through the price experiment. This time, the identification strategy is essentially comparing β_{st} between songs whose $s < t$ (already treated) and songs whose $s > t$ (not yet treated). Again, since β_{st}'s have a similar trend with t across s, the sales of singles whose prices were changed at different time points trended similarly over time until the price changes. Therefore, "not yet treated" songs themselves can also be a good control group, and we can reliably identify the price effect by comparing the "treated" and "not yet treated" groups. Given this, in the next section we run a reduced-form regression based on the samples in Figure 3, (a) and (b), and confirm that both elasticity estimates are significant and very similar.

4. Reduced-Form Analysis

Before we present our structural model, we first conduct a series of reduced-form analyses to explore data patterns. We compare these reduced-form estimates to our structural results in later sections.

We use the $0.30 "exogenous" price increase described earlier as a quasi-experiment to identify the own- and cross-price elasticities in our experiment. Because we expect digital song sales to change over time regardless of price, we use tracks that remained at $0.99 as a control group, asking how sales for $1.29 priced songs change over and above any change for the control group. As discussed above, the identifying assumption behind this difference-in-difference model is that $0.99 songs provide a counterfactual that indicates how sales of $1.29 priced songs would have changed in the absence of the price increase (see discussion in §3).

Formally, let a index albums and t index the time period in our data. We rank songs in an album based on their average sales prior to the price experiment and use j to denote the rank of the songs in the album and J_a to denote the total number of songs in the album a. A rank is assigned to a specific song and is carried through the whole period. In the following regression, dummy variable I_{ajt} is set to 1 as long as song j is priced at $1.29 in week t, and its coefficient κ measures the effect of the $0.30 price change. Changes to other songs' prices may affect the sales of the focal song, and as such, in Equation (1), we also include a dummy variable I_{a-jt}, which is set to 1 when the price of at least one other song in the same album is increased to $1.29 to control for possible cross-price elasticity between single songs. We also include song-specific dummy variables (c_{aj}) to control for song-level unobserved heterogeneity and week-specific dummy variables (D_t) to control for the any common weekly market-level shocks that may have affected all singles:

$$LogSales_{ajt} = c_{aj} + \kappa I_{ajt} + \rho I_{a-jt} + \sum_t d_t D_t + \xi_{ajt}. \quad (1)$$

The estimation results for Equation (1), displayed in Table 5, show an estimate of −0.145 for κ, indicating that when the single song's price changes from $0.99 to $1.29 (a 30% increase), sales decrease by 14.5%. Thus, under the assumption of constant price elasticity, our results suggest that a 1% increase in single song's price will lead to a 0.48% drop in sales. The estimate shows that the demand of digital music was relatively inelastic at preexperimental prices, consistent with Belinfante and Johnson's (1982) argument. In addition, the coefficient of I_{a-jt} is negative and significant, providing evidence that other songs' price increases negatively

Table 5 Reduced Form—Regression of Song Log Sales on Own- and Cross-Price Changes; Common Own-Price Elasticity for Songs with Different Ranks (Full Sample)

Variable	Coefficient	Standard error[a]
I_{ajt}	−0.145***	0.014
I_{a-jt}	−0.046***	0.008
Song fixed effect	Yes	
Week dummies	Yes	
$n = 3{,}709$, $T = 7$–30,[b] $N = 108{,}257$; adjusted R-squared: 0.072		

[a]Robust standard errors controlling for heteroskedastic errors across songs and autocorrelation within the same song are reported.

[b]"T" in the table represents the number of observations for a particular song. "$T = 7$–30" means that the range of the number of weekly observations of these songs is from 7 to 30. In our data set, most songs (and their parent albums) have 30 observations, whereas a very small number of songs have fewer observations (less than 5%).

***$p < 0.001$.

Table 6 Reduced Form—Regression of Song Log Sales on Own- and Cross-Price Changes; Common Own-Price Elasticity for Songs with Different Ranks (Including Only Songs That Eventually Underwent a Price Change)

Variable	Coefficient	Standard error[a]
I_{ajt}	−0.124***	0.019
I_{a-jt}	−0.066***	0.019
Song fixed effect	Yes	
Week dummies	Yes	
$n = 1{,}457$, $T = 7$–30, $N = 42{,}887$; adjusted R-squared: 0.119		

[a]Robust standard errors controlling for heteroskedastic errors across songs and autocorrelation within the same song are reported.
*** $p < 0.001$.

Table 7 Reduced Form—Regression of Song Log Sales on Own- and Cross-Price Changes; Heterogeneous Own-Price Elasticity for Songs with Different Ranks (Full Sample)

Variable	Coefficient	Standard error[a]	Variable	Coefficient	Standard error[a]
$I_{a1t} \times Rank\ 1$	−0.130***	0.029	$I_{a11t} \times Rank\ 11$	−0.161*	0.070
$I_{a2t} \times Rank\ 2$	−0.072***	0.021	$I_{a12t} \times Rank\ 12$	−0.150+	0.090
$I_{a3t} \times Rank\ 3$	−0.160***	0.033	$I_{a13t} \times Rank\ 13$	−0.251*	0.110
$I_{a4t} \times Rank\ 4$	−0.082*	0.034	$I_{a14t} \times Rank\ 14$	−0.207*	0.087
$I_{a5t} \times Rank\ 5$	−0.133***	0.040	$I_{a15t} \times Rank\ 15$	−0.373***	0.011
$I_{a6t} \times Rank\ 6$	−0.206***	0.046	$I_{a16t} \times Rank\ 16$	−0.180***	0.011
$I_{a7t} \times Rank\ 7$	−0.176***	0.048	$I_{a17t} \times Rank\ 17$	−0.138***	0.010
$I_{a8t} \times Rank\ 8$	−0.264***	0.051	$I_{a18t} \times Rank\ 18$	0.120[b]	0.195
$I_{a9t} \times Rank\ 9$	−0.217***	0.063	$I_{a19t} \times Rank\ 19$	—	—
$I_{a10t} \times Rank\ 10$	−0.270***	0.081	I_{a-jt}	−0.048***	0.008
Song fixed effect	Yes		Week dummies	Yes	
$n = 3{,}709$, $T = 7$–30, $N = 108{,}257$; adjusted R-squared: 0.046					

[a]Robust standard errors controlling for heteroskedastic errors across songs and autocorrelation within the same song are reported.
[b]There is only one rank 18 and no rank 19 songs that experience the price change. Therefore, the coefficient of $I_{a18t} \times Rank\ 18$ is not significant, whereas the coefficient of $I_{a19t} \times Rank\ 19$ cannot be estimated. The estimated effect of other songs' price change is still negative and significant.
$^{+}p < 0.1$; $^{*}p < 0.05$; $^{***}p < 0.001$.

affect focal song sales, possibly because of consumers switching from single songs to the album. Therefore, when we explore digital music pricing strategies, we must consider song–album interdependence.

In Table 6, we run the same regression as in (1), but this time using only those songs that eventually underwent a price change. These estimation results again indicate that increases in the price of the focal song and in the price of other songs in the same album have a negative impact on focal songs sales. Compared with the results in Table 5, the point estimate of the coefficient of I_{ajt} is smaller in magnitude, and the point estimate of the coefficient of I_{a-jt} is larger in magnitude, which is consistent with what we observe in Figure 3, but the differences are not significant. This suggests that either nonexperimented songs or experimented songs that had not yet undergone a price shift are suitable controls for experimented songs that had undergone the price shift.[5]

In addition, it is also likely that songs with different rankings may have different price elasticities. In Equation (2) we interact the rank dummy variables (R_j, within-album sales rank) with the price change dummy variables (I_{ajt}):

$$LogSales_{ajt} = c_{aj} + \sum_{j=1}^{J_a} \kappa_j I_{ajt} \times R_j + \rho I_{a-jt} + \sum_t d_t D_t + \xi_{ajt}. \quad (2)$$

The results (shown in Table 7) suggest a trend where higher-ranked songs are less affected by their own price change than lower-ranked songs are. We note that in the reduced-form estimations, the effect of the concurrent price changes happening to other songs in the album is captured by a single term I_{a-jt}, and the album price change is not taken into account. Therefore, the estimates of price elasticity are noisy.

In Equation (3) we further explore how songs' price changes will affect album sales. The dependent variable in this equation is the log of album sales in each week. The independent variable of interest is I_{ast}, a dummy variable that takes the value 1 when the price of at least one (and possibly more than one) of the single songs in album a is increased to \$1.29. In the regression we include both album and time dummies. Although we are aware that album price (p_{aAt}) could be endogenous, we ignore this potential endogeneity for now. Below, we discuss how we deal with this problem in our structural model:

$$LogSales_{aAt} = c_{aA} + \kappa_A \mathrm{Log}(p_{aAt}) + \rho I_{ast} + \sum_t d_{At} D_t + \xi_{aAt}. \quad (3)$$

As shown in Table 8, the coefficient on I_{ast} is positive and significant, providing some evidence that after we control for album price (p_{aAt}) changes, the increase in single song prices will lead to a significant increase in album sales as predicted by bundling theory.[6] However, as noted above, $I_{ast} = 1$ when one or more of the songs in the album changed prices, and in many cases in our data, multiple songs in a single album changed prices. As such, one should not interpret I_{ast} as a

[5] We even further divide all songs that underwent the price change into four groups based on the week in which a song's price was increased. Songs whose price change occurred from week 17–19 form group 1, those whose price change occurred from week 20–22 form group 2, those whose price change occur from week 23–26 form group 3, and those whose price change occurred from week 27–30 form group 4. The rationale behind this is that songs in each of the groups have even more similar sales trends, as shown in Table 4. The estimates of the coefficient of I_{ajt} are still similar in each of the four groups.

[6] Note that this cross-price elasticity refers to digital albums, not physical CDs.

Table 8 Reduced Form—Regression of Album Log Sales on Album Price and Song Price Changes

Variable	Coefficient	Standard error[a]
$\text{Log}(p_{aAt})$	-1.785^{***}	0.231
I_{ast}	0.100^{*}	0.041
Album fixed effect	Yes	
Week dummies	Yes	
$n = 286$, $T = 7$–30, $N = 8{,}342$; adjusted R-squared: 0.069		

[a]Robust standard errors controlling for heteroskedastic errors across albums and autocorrelation within the same album are reported.
$^{*}p < 0.05$; $^{***}p < 0.001$.

price elasticity from a \$0.30 price change. Rather, we are interested in the coefficient on I_{ast} primarily to understand whether there is a correlation between song price change and album sales as a way of guiding our structural modeling below.

5. Structural Model

In §4, we provided evidence in our data that the changes in single song prices will affect not only the sales of the songs experiencing the price change, but also sales of other songs in the same album and sales of the album itself. However, the mechanism driving these effects is complicated, and the reduced-form model cannot disentangle these different effects readily. For example, in many cases, prices of multiple songs increase simultaneously; therefore, it is hard to tease apart the net effect of each price change. In addition, the reduced form does not provide any guidelines on the customer decision-making framework, an important criterion for doing any counterfactuals. Since one of our research goals is to provide recommendations for optimal single and album prices, we need a formal model. Our empirical specification is based on an underlying model of individual consumer utility maximization. Users can either choose individual song(s) or can choose to buy the entire album (a bundle of songs). Thus the singles (individual components) and the album are potential substitutes for each other. In other words, when setting single song prices and the price for an album, economic theory would predict that Label X acts as a monopolist, aiming to maximize the joint profit from the sales of both the full-length album and the singles. We note that the songs are not directly competing with each other. In the following utility function, subscripts a, i, t, and j are defined in the same way as in the previous section. In each period, consumers' indirect utility from purchasing the jth song in album a can be written as

$$U_{iajt} = \theta_{aj} + \alpha p_0 + \gamma I_{ajt} + \sum_t d_t D_t + \mu_{ajt} + \varepsilon_{iajt},$$
$$\text{where } \varepsilon_{iajt} \sim MVN(0, I_{J_a}), \quad (4)$$

$$\delta_{ajt} = \chi_{aj} + \gamma I_{ajt} + \sum_t d_t D_t + \mu_{ajt}, \quad (5)$$

where δ_{ajt} is the mean utility of the jth song in album a in period t across individuals, ε_{aijt} is the idiosyncratic error term, and the mean utility δ_{ajt} is composed of four parts: χ_{aj}, which is the combination of the constant term in the utility of the jth song in album a, denoted as θ_{aj}, and the price component αp_0, where $p_0 = 0.99$; γI_{ajt}, which is the combination of I_{ajt}, a dummy variable indicating songs with prices of \$1.29, and γ, which captures the mean utility decrease when the price of a song is raised from \$0.99 to \$1.29; $\sum_t d_t D_t$, which is the combination of D_t, a set of week dummies, and d_t, which measures aggregate demand shifts in the digital music market in each week; and μ_{ajt}, which is an unobserved market-level random disturbance for the jth song in album a in period t. We note that μ_{ajt} is assumed to be zero mean and is uncorrelated with I_{ajt} and the time dummies, but we allow μ_{ajt} to be heteroskedastic across songs and correlated within the same song. Notice μ_{ajt} affects all users in the same way. This is unlike ε_{iajt}, the individual-specific random shock, assumed to follow an independent and identically distributed (i.i.d.) standard normal distribution. Consumer heterogeneity enters the model only through the separable additive random shocks, ε_{iajt}.[7] Note that in the structural model, we normalize the utility consumers derive from consuming "outside goods" to 0. Music available in other distribution channels and digital music available at other online outlets are all included in the composite "outside good." Therefore, the utility of purchasing digital music we model here is relative to the utility that consumers derive from consuming outside goods. Similarly, we assume that the utility consumers derive from purchasing an album is given by

$$U_{iaAt} = \theta_{aA} + \alpha p_{aAt} + \sum_t d_{At} D_t + \mu_{aAt} + \sum_{j=1}^{J_a} \varepsilon_{iajt}, \quad (6)$$

$$\delta_{aAt} = \theta_{aA} + \alpha p_{aAt} + \sum_t d_{At} D_t + \mu_{aAt}, \quad (7)$$

where δ_{aAt} is the mean utility consumers can get from purchasing the album. We assume that every potential consumer's individual-specific shock on album utility is the sum of the individual's random shocks on all songs in the album; that is, if one consumer likes some songs in the album very much, she will also value the album higher and vice versa. In Equation (6), θ_{aA} is the average utility a consumer can derive from purchasing the album at price zero; μ_{aAt} is the album-level aggregate shock in period t, which affects all consumers' utility in the same way; and αp_{aAt} measures

[7] We note that one could also include song attributes in the utility function. However, in our setting, almost all attributes observed in the data, such as genre and artist, are time-invariant attributes, which are absorbed in the song-specific constant θ_{aj}.

the effect of the album price (p_{aAt}) on consumers' utility. Different from the prior bundling literature in which the value of the album is assumed to be the sum of the values of its components (some of the bundling papers assume free disposal), we assume an independent value of the album δ_{aAt},[8] which connects consumers' valuation on single songs and the album through an individual-specific random shock ε_{iaAt}. This is because if the digital album is simply a bundle of songs and consumers do not receive additional value when purchasing the album compared to purchasing all songs in the album separately, consumers will never purchase the album when the album price is higher than the sum of all single songs' prices. However, in our data set, we frequently observe that consumers purchase the album, even when the album price is higher than the sum of all single songs' prices. Our flexible specification of album value allows for this possibility.

In each period, consumer i will consume song j in album a as long as $U_{iajt} > 0$, or, equivalently, when

$$\varepsilon_{iajt} > -\left(\chi_{aj} + \gamma I_{ajt} + \sum_t d_t D_t + \mu_{ajt}\right). \tag{8}$$

The probability that (8) holds is

$$P(U_{iajt} > 0) = P\left(\varepsilon_{iajt} > -\left(\chi_{aj} + \gamma I_{ajt} + \sum_t d_t D_t + \mu_{ajt}\right)\right) = \Phi\left(\chi_{aj} + \gamma I_{ajt} + \sum_t d_t D_t + \mu_{ajt}\right).$$

However, this is not the probability that song j is purchased. To consume song j, consumers have another choice: to buy the whole album. Consumers will choose to purchase the album instead of single song(s) when the utility they can derive from purchasing the whole album exceeds the sum of the utilities they get from purchasing the songs separately. We follow the common assumptions in the bundling literature that there are no demand-side complementarities from consuming particular songs together.[9] Under this assumption, the utility consumers receive when they choose to purchase multiple singles is simply the sum of the utilities they obtain from consuming each of the singles.

[8] However, we do not allow for heterogeneity in consumers' preference toward the album (or the single); i.e., in our model, we do not separately distinguish consumers who prefer getting the full album and consumers who prefer to cherry pick hit songs, all other things being equal. Since we only have aggregate sales data and limited price variation in the data, we cannot identify different mean valuations of an album (θ_{aA}) for different types of consumers.

[9] There are no demand-side complementarities unless the entire album is purchased, at which point the fact that we allow album value to be higher than the sum of its components could reflect complementarities across songs.

The purchasing probabilities of the album and single songs can then be expressed as

$$s_{aAt} = P(\text{buying album } a) = P\left(0 < \sum_{j=1}^{J_a} U_{iajt} I\ (U_{iajt} > 0) < U_{iaAt}\right), \tag{9}$$

$$s_{ajt} = P(\text{buying song } j \text{ in album } a) = P\left(U_{iajt} > 0 \cap \sum_{j=1}^{J_a} U_{iajt} I\ (U_{iaAt} > 0) > U_{iaAt} > 0\right), \tag{10}$$

where $I(U_{iajt} > 0) = 1$ if $U_{iajt} > 0$, and $I(U_{iajt} > 0) = 0$ otherwise.

Below, we summarize the key assumptions of this model:

(1) Each album is a monopolistic market.

(2) We assume that within an album, consumers' purchase decisions on single songs in the album are independent, and their choice between purchasing multiple single songs and the album is exclusive.

(3) The price parameters are the same in both an individual's utility function for purchasing tracks (Equation (4)) and their utility function for purchasing the album (Equation (6)).

(4) ε_{iajt} follows an i.i.d. standard normal distribution (i.e., no correlation among songs; however, we relax this assumption later in robustness checks). But the consumer's individual-specific shock on the album utility is the sum of the individual's random shocks on all songs in the album, i.e., $\varepsilon_{iaAt} = \sum_{j=1}^{J_a} \varepsilon_{iajt}$.

(5) Consumer taste and price sensitivity parameters are homogeneous, and consumer heterogeneity enters the utility function only through the separable additive random shocks, ε_{iajt}.

In this analysis we consider each album as a monopolistic market, because in the music industry albums are highly differentiated, and so there is little substitution across albums. We further assume that the price sensitivity parameter is the same in both the song utility function and in the album utility function. The price sensitivity parameter captures the opportunity cost of spending \$1 on music consumption versus spending this \$1 on the outside good, and therefore should be the same for either a single song purchase or an album purchase. Notice that we do allow for differences in the valuation of an album versus the sum of the valuations of all songs in the album. We assume that a consumer's individual-specific shock on the album utility is the sum of the individual's random shocks on all songs in the album, to reflect the fact that if a consumer has high valuation for individual songs in an album, she will also value the album higher and vice versa. This assumption also makes individual-specific valuation correlated across songs and albums.

The last assumption we make is that the consumers have both a common taste parameter and a common price sensitivity parameter. Below, we discuss the implications of homogeneous consumer taste and price sensitivity:

1. *Homogeneous valuation*. In our data, all song attributes, such as genre, artist, or whether it is a catalog/new release, are time invariant. As such, consumers' preferences of song attributes are all in the fixed effect terms θ_{aj}.[10] In other words, $\theta_{aj} = \sum_{k=1}^{K} \beta_k X_k$, where X_k represents the kth preference attribute, and β_k represents consumers' mean marginal utilities associated with this attribute. Now let us assume that consumers are heterogeneous in their preferences. For an individual i, her preference for attribute k, denoted as β_{ik}, can be decomposed as $\beta_k + \tilde{\beta}_{ik}$, in which $\tilde{\beta}_{ik}$ represents the deviation of this individual's preference from the mean level and is assumed to follow $N(0, \sigma_k^2)$.[11] With this, Equation (4) can be rewritten as

$$U_{iajt} = \theta_{aj} + \alpha p_0 + \gamma I_{ajt} + \sum_t d_t D_t + \mu_{ajt} + \sum_{k=1}^{K} \tilde{\beta}_{ik} X_k + \varepsilon_{iajt}.$$

But notice that X_k is same whether the song is chosen as an individual song or it is chosen within the album because they have identical attributes.[12] In fact, in our data, the attributes of all songs in the same album are the same, and $\sum_{k=1}^{K} \tilde{\beta}_{ik} X_k$ is common across all songs for the same individual i and is time invariant. In other words, for each individual i, $U_{iajt} \sim N(\delta_{ajt} + d_i, 1)$, where d_i is $\sum_{k=1}^{K} \tilde{\beta}_{ik} X_k$, which is fixed for all j's, and itself follows $d_i \sim N(0, \sigma_d^2)$ across different i's. We let $\eta_{iajt} = \sum_{k=1}^{K} \tilde{\beta}_{ik} X_k + \varepsilon_{iajt}$, where at the aggregate level, the distribution of η_{iajt} is $N(d_i, 1) \mid d_i$ and $d_i \sim N(0, \sigma_d^2)$. This is equivalent to assuming η_{iajt} to be correlated across different songs (j) and independent across individuals (i), with identical off-diagonal elements in the variance–covariance matrix. This situation is later examined in the Robustness Check section.

2. *Homogeneous price sensitivity*. This situation is somewhat similar to the situation we discussed above where $U_{iajt} \sim N(\delta_{ajt} + d_{ijt}, 1)$ and where d_{ijt} is $\tilde{\alpha}_i p_{ajt}$, which follows $N(0, p_{ajt}^2 \sigma_\alpha^2)$. Here, $\tilde{\alpha}_i$ represents the deviation of individual i's price sensitivity from the mean level, and we again define $\eta_{iajt} = \tilde{\alpha}_i p_{ajt} + \varepsilon_{iajt}$. Notice that here we add subscripts j and t into d_{ijt}, because d_{ijt} is a function of p_{ajt}, which is time-variant and song specific. When the price of a song increases, the most obvious effect is that δ_{ajt} drops. It also has an effect on the variance and covariance matrix of η_{iajt}: the variance of d_{ijt} increases, increasing elements in the jth column and jth row of the variance–covariance matrix of η_{iajt}. Thus, price changes have a complex effect on aggregate sales, and with limited price variation (only two price points) in our data, it is very difficult to tease apart its effects on δ_{ajt} and on the variance–covariance matrix.

Note that although we impose assumptions on individuals' utility functions, the elasticities of songs (of different ranks) and the album are still estimated from the data, because the key parameters in the utility functions are recovered empirically from our experimental data. In the next section, we will discuss how we recover these parameters.

6. Estimation

6.1. Estimation Procedure

To estimate the choice model above using aggregate data, we have to combine the well-known BLP method (Berry et al. 1995) with a similar estimation method designed for Probit models introduced by Chintagunta (2001). The principle underlying our estimation is simple: Obtain estimates of δ_{ajt} and δ_{aAt} by equating S_{ajt} (and S_{aAt}), the observed proportions of potential consumers who purchase song j (and album a) in period t, to s_{ajt} (and s_{aAt}), the probabilities predicted by the model as calculated from Equations (9) and (10). Given the market size (the number of the potential consumers in the market), we can calculate S_{ajt} (and S_{aAt}) by dividing sales of the album and each single song by the market size.

The estimation procedure involves two steps:

Step 1: *Estimate* δ_{ajt} (*including* δ_{aAt}). For any given album a and period t:

(1) Make initial guesses for the parameters δ_{ajt} and δ_{aAt}.

(2) The "inner loop" computation takes place. Let M denote market size (we will discuss how we select M later in this section), and let J_a denote the number of singles in the album. For each consumer (denoted by i), do the following:

(i) For each song j, draw a ε_{iajt} for all J_a singles from $MVN(0, I_{J_a})$.

[10] There could be time-variant factors, such as promotional activities, media exposure, and so on, affecting single song sales. However, most of songs in our sample are of older vintage. The relative popularity of these songs has been revealed, and there is little promotional activity for these albums. Despite a decline trend, songs in these albums exhibit stable sales (e.g., few songs suddenly become hits). The within-album ranks are also stable. Therefore, we do not incorporate time-variant attributes in the model, and X_k only includes time-invariant characteristics. Because of this, we will not be able to study the interplay between marketing and promotional strategy and the pricing strategy in this paper. The precise timing of price increase and decrease is not something we can analyze with our data either.

[11] We can further allow for nonzero correlation among an individual's preferences for different attributes, i.e., allow vector $\tilde{\beta}_i = (\tilde{\beta}_{i1}, \tilde{\beta}_{i2}, \ldots, \tilde{\beta}_{iK})$ to follow a multivariate normal distribution $\tilde{\beta}_i \sim MNV(0, \Sigma)$. In this case, $d_i \sim N(0, X'\Sigma X)$, where X is an $K \times 1$ vector of song attributes.

[12] This is unlike BLP, where all products are differentiated especially along the X dimension.

(ii) Calculate the utility this consumer can generate from purchasing the most popular single U_{iajt}, which is $\delta_{ajt} + \varepsilon_{iajt}$ where $j = 1$.

(iii) Repeat (ii) above for all other singles $j = 2$ to J_a.

(iv) Calculate the utility consumer i can generate from purchasing the whole album using $U_{iaAt} = \delta_{Ajt} + \sum_{j=1}^{J_a} \varepsilon_{iajt}$.

(v) Find out the singles (j's) that provides this consumer with positive utility, i.e., find for which j's U_{iajt} is greater than 0.

(vi) Calculate $\sum_{j=1}^{J_a} U_{iajt} I\ (U_{iajt} > 0)$, the total utility consumer i can generate from purchasing the singles that provide positive utility (the ones that are selected in step (v)).

(vii) Compare U_{iaAt} and $\sum_{j=1}^{J_a} U_{iajt} I\ (U_{iajt} > 0)$. If $U_{iaAt} > \sum_{j=1}^{J_a} U_{iajt} I\ (U_{iajt} > 0)$, this consumer will purchase the album. If $U_{iaAt} < \sum_{j=1}^{J_a} U_{iajt} I\ (U_{iajt} > 0)$ and $\sum_{j=1}^{J_a} U_{iajt} I\ (U_{iajt} > 0) > 0$ (at least one of the singles provides the consumer with positive utility), then this consumer will purchase the singles (j's) that satisfy $U_{iajt} > 0$. Otherwise, this consumer will purchase nothing.

(viii) Repeat steps (i)–(vii) for all M potential consumers in the market. Since we have simulated each consumer's choices, we can calculate the simulated sales. Simulated market shares $S_{a1t}, S_{a2t}, \ldots, S_{aJ_at}$, and S_{aAt} can then be obtained by dividing the simulated sales by M.

(3) Plug the simulated market share into the objective function $\|S_{at} - s_{at}\|^2$. In the outer loop, search for the set of δ_{ajt} and δ_{aAt} that can minimize the objective function, i.e., $\hat{\delta}_{at} = \arg\min_{\delta_{at}} \|S_{at} - s_{at}\|^2$.

To ensure the precision of $\hat{\delta}_{at}$, we repeat the minimization procedure 100 times and select the set of $\hat{\delta}_{at}$ that gives the smallest value of $||S_{at} - s_{at}||^2$.

Step 2: *Estimate parameters in the utility function*. Since μ_{ajt} is a random shock (i.e., independent of I_{ajt} and D_t), χ_{aj} and γ can be estimated with linear regressions of song-level fixed effects. Notice μ_{ajt} will also pick up possible numerical errors in $\hat{\delta}_{at}$ produced in the first step of the estimation procedure.

One additional assumption we need to make to estimate the model is about the market size M. The selection of the market size can affect the resulting estimates. One of the important rules for selecting M is that it should give a reasonable share of no-purchase incidences. In BLP, market size is defined as the number of households in the economy, whereas Nevo (2000) sets market size for ready-to-eat cereal to be one serving of cereal per capita per day. In the same spirit, one straightforward definition of M in our context is to set market size equal to the total number of users of Retailer Y's online store. Although this definition of M can give us a reasonable share of no-purchase incidences, it will also lead to extremely small market shares for many albums, because the average sales of many singles and albums are less than 100. The extremely small share values will yield imprecise estimates of δ_{ajt}, and therefore the total number of users of the online store is not a proper definition of M. In this paper, the size of the market is defined as twice the maximum sales observed in the album before the experiment started (week 17). We define market size M this way for two reasons: (1) This definition allows a significant share of nonpurchasing choices while keeping the share of songs sufficiently large to ensure the precision of the estimates of δ_{ajt}, and (2) defining M to be proportional to the maximum sales performs a normalization function, which allows us to combine multiple albums with different absolute sales to estimate a single price sensitivity parameter and a common set of intrinsic values of songs in a typical album.

6.2. Identification

In Equation (5), it is clear that any time-invariant attribute, such as genre or catalog/new release, will be absorbed in the song-specific intercept χ_{aj}. The only time-variant attributes are the price changes, which we denote as I_{ajt}, and the weekly dummies. We argue that a song-specific fixed-effects regression can be used here to identify the main parameter of interest, γ. Our reduced-form regressions have already established this identification. The intuition is that after controlling for song-specific fixed effects and the time fixed effects for a typical song, we compute the difference in the mean values between songs that experience the price change in each period and those that do not experience the price change in the same period. The estimate of γ is obtained by averaging the differences. Since the price increase is \$0.30, the estimate of the price sensitivity parameter is $\gamma/0.3$, as price enters the consumers' utility functions linearly.

The song-specific intercept χ_{aj} (the fixed effect dummy in (5)) can be broken down further to gain more insights about the value distribution of songs in an album. Specifically, χ_{aj} can be expressed as

$$\chi_{aj} = \sigma_a + \sum_{j=2}^{J_a} \varphi_j R_j + \zeta_{aj}. \qquad (11)$$

In Equation (11), σ_a is the album-level fixed effect capturing the value of the most popular song in album a, R_j is the rank indicator, and φ_j measures rank-specific effects. We also assume the last term, ζ_{aj}, is random, such that φ_j reflects the relative intrinsic value difference between the jth ranked song and the first ranked song in the same album.

6.3. Estimation Results

6.3.1. Price Parameter. We follow the two-step procedure outlined above to estimate our model, and the results of Equations (5) and (11) are displayed in Tables 9 and 10, respectively. In Table 9, we see that the estimate of γ is -0.119,[13] which implies that the price sensitivity parameter is $\alpha = -(0.119/0.3) = -0.397$. This means that a \$1 increase in price leads to 0.397 drop in utility.

6.3.2. Rank Analysis. Table 10 summarizes the estimates (from Equation (11)) of the relative song value by rank with respect to the value of the highest-ranked song (i.e., the first ranked song). This result is shown graphically in Figure 4. The estimates show that there is significant heterogeneity in the intrinsic values of songs with different ranks. For example, the second ranked song is valued 0.629 lower relative to the top ranked song. However, the intrinsic value decreases faster among higher-ranked songs, and the difference is smaller among the lower-ranked songs. Moreover, the decreasing trend in song intrinsic value is consistent with the shape of the average weekly sales by rank we show in Figure 1.

6.3.3. Album Value Analysis. Since in our first estimation step we also obtain the estimate of δ_{aAt}, denoted as $\hat{\delta}_{aAt}$, the mean utility level of albums, we can explore the relationship between single songs and albums explicitly. To do this, we define "album net value" as

$$Y_{at} = \hat{\delta}_{aAt} - \hat{\alpha} p_{aAt}, \qquad (12)$$

which captures the utility consumers receive excluding price components (i.e., the net value of the album). Note that to recover Y_{at}, we rely on the assumption that the price sensitivity parameter α is the same for singles and the album. We further explore how the album value evolves over time by estimating the following equation:

$$Y_{at} = \pi_a + \sum_t \lambda_t D_t + v_{at}. \qquad (13)$$

In Equation (13), π_a is the album fixed effect, and v_{at} is the random disturbance. The estimate for the mean of π_a is -0.752, meaning that when the album price is zero, the average album value is -0.752. This should not be surprising because this is the unconditional valuation of the album. In other words, because of many nonbuyers, the average valuation must be a small number.

[13] This estimate of this parameter is robust when we include log(*weeks since release*) variable in the regression or eliminate albums whose prices ever changed during our study period (in fact, for the vast majority of albums in our data, album price remained the same). This reflects the fact that possible omitted variables, log(*weeks since release*) and album prices, are not a big concern, given the nature of our experimental data.

Table 9 Song-Level Fixed-Effect Regression of Song Mean Utilities

Variable	Coefficient	Standard error[a]
l_{ajt}	-0.119^{***}	0.013
Song fixed effect	Yes	
Week dummies	Yes	

$n = 3{,}709$, $T = 7$–30, $N = 108{,}257$; adjusted R-squared: 0.015

[a] Robust standard errors controlling for heteroskedastic errors across songs and autocorrelation within the same song are reported.
$^{***}p < 0.001$.

Table 10 Album-Level Fixed-Effect Regression of the Estimated Song-Specific Intercept on Rank Dummies

Variable	Coefficient	Standard error[a]	Variable	Coefficient	Standard error[a]
Rank 2	-0.629^{***}	0.034	*Rank* 11	-1.846^{***}	0.032
Rank 3	-1.057^{***}	0.038	*Rank* 12	-1.887^{***}	0.034
Rank 4	-1.307^{***}	0.039	*Rank* 13	-1.911^{***}	0.033
Rank 5	-1.477^{***}	0.037	*Rank* 14	-1.944^{***}	0.035
Rank 6	-1.588^{***}	0.039	*Rank* 15	-1.981^{***}	0.035
Rank 7	-1.670^{***}	0.034	*Rank* 16	-2.023^{***}	0.037
Rank 8	-1.714^{***}	0.033	*Rank* 17	-2.029^{***}	0.041
Rank 9	-1.772^{***}	0.033	*Rank* 18	-2.065^{***}	0.041
Rank 10	-1.812^{***}	0.033	*Rank* 19	-2.121^{***}	0.036

[a] Robust standard errors controlling for heteroskedastic errors across albums and autocorrelation within the same album reported.
$^{***}p < 0.001$.

Another point worth mentioning is that the estimate of π_a is larger than the sum of θ_{aj}'s, indicating that the value of the album exceeds the sum of the values of the songs in the album. One possible explanation is that individuals associate a cost with the effort necessary to separately purchase multiple single songs versus a single album. Since individuals incur this cost whenever they purchase a single or an album, we

Figure 4 Coefficients of Rank Dummies

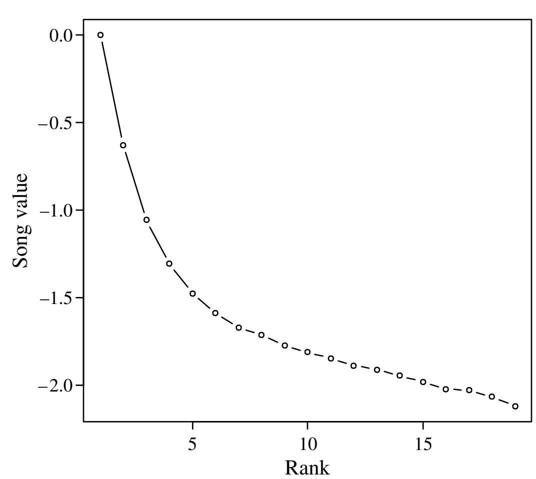

Notes. The first point represents the first ranked song. Because we use the first ranked song as baseline, its relative value is zero.

cannot separate it out from the song and album values. If that is the case, the song (album) value we estimated here is actually a combination of the net value of the song (album) and the cost of the effort necessary to purchase the song or album.[14] Individuals incur higher total hassle cost when purchasing singles separately, as compared to purchasing the album at once.

6.4. Price Elasticity Analysis

In the previous section we estimated the price change effect, γ. From the estimate of γ, we then recovered the price sensitivity parameter α in the utility function. We can now use these estimates to recover the own- and cross-price elasticity of demand. In our structural analysis, the probabilities that an album and the singles within the album will be purchased do not have a closed-form solution. Therefore, we have to use numerical methods to approximate the price elasticity. Note that in our Probit choice model, elasticity is not constant at different price points, whereas in the reduced-form models, the elasticity is assumed to be constant. To facilitate comparison between the reduced-form results and the structural results, we simulate a price change from \$0.99 to \$1.29, which is the same as what we observe in the data. We then present the predicted effect of the price change based on the reduced-form and structural specifications. Strictly speaking, we are comparing the percentage change in sales after the 30% price increase is executed, which we call the "price effect." Although this is not the same as the formal definition of elasticity, the nature of the price effect approximates elasticity, and we use the two terms interchangeably.

Consider a typical album with 13 songs (the median of the numbers of songs in our sample). We use the estimates in Table 10 to calculate the values of different songs in the album. Note that Table 10 shows the values of songs of different rankings relative to the value of the most popular song, which is captured in the album-specific intercept. We take the average of all the album-specific fixed effect coefficients for 13-song albums as χ_{a1} and add the first 12 coefficients in Table 10 to the χ_{a1} to obtain χ_{aj} for $j = 2$–13. Furthermore, we exclude the price component $\theta_{aj} = \chi_{aj} - \alpha p_0$. In the simulation, we also ignore the effect of seasonality. We first recover the rank-specific song value χ_{aj} including αp_0 for $j = 1$–13, and then use the estimate of the price coefficient to recover the pure song value θ_{aj} (shown in Table 11).

[14] The effort necessary to take actions online has been studied widely in the search cost literature. For example, Hann and Terwiesch (2003) estimate that the cost to rebid in a name-your-own-price auction is between \$3 and \$6, Bajari and Hortaçsu (2003) estimate that the cost of entering an eBay auction is \$3.20, and Brynjolfsson et al. (2010) estimate that the cost of searching on a shopbot is between \$6 and \$7.

Table 11 Song Value Distribution

j	χ_{aj}	θ_{aj}	j	χ_{aj}	θ_{aj}
1	−0.653	−0.260	8	−2.367	−1.974
2	−1.282	−0.889	9	−2.425	−2.032
3	−1.710	−1.317	10	−2.465	−2.072
4	−1.960	−1.567	11	−2.499	−2.106
5	−2.130	−1.737	12	−2.540	−2.147
6	−2.241	−1.848	13	−2.564	−2.171
7	−2.323	−1.930			

Then, we simulate sales of all 13 songs and the sales of the album when the price of each of the songs in the album and the price of the album increase by 30%, respectively. We define the market size as 5,000 and assume that all potential consumers in the market are homogeneous and independent. The assumption on market size will have no impact on the percentage change in sales. We first simulate the sales volume under the current price, and then change the price of a song or the album at a time and ask how the sales of all songs and the album change. For each price change, we run 200 iterations of the simulation and report the rounded average sales numbers over the 200 iterations. Two hundred iterations are sufficient because the difference of average sales across any 200 iterations is less than or equal to 2. The simulation results for a typical album with 13 songs are displayed in Table 12.

In Table 12, each row represents a case where the price of a song (or the album) is raised, whereas each column indicates the sales of each song (or the album). For example, the "baseline" row shows the sales of singles and the album when the prices of all singles are \$0.99 and the price of the album is \$9.99. The numbers in row 1 indicate the sales of all songs and the album when the price of the first song is raised to \$1.29 (and the other songs and the album have the same price as they do in the baseline case). The percentage change is relative to sales in the baseline case. Therefore, the diagonal elements show the own-price effect, whereas the off-diagonal elements show the cross-price effect. The first number in each cell is the simulated sales volume, and the second number is the percentage change. Since the price change is 30% and the price elasticity is not constant, we can only say that the price elasticity approximates the percentage change in sales reported in each sale divided by 30. The simulation error is around 2, suggesting that an absolute change in sales greater than 4 is roughly statistically significant. The numbers in bold in Table 12 would be considered significant.

From the simulation results, we can see that the own-price effect is more significant for songs with lower ranks than for songs with higher ranks: Consumers are much more inelastic with respect to the price of

Table 12 Own- and Cross-Price Effects for Single Songs and Album

	1	2	3	4	5	6	7	8	9	10	11	12	13	Album
Baseline	1,150	434	181	102	68	50	40	35	30	27	25	21	21	359
1	**974** **−15.30%**	**429** **−1.15%**	182 0.55%	103 0.98%	68 0.00%	49 −2.00%	39 −2.50%	36 2.86%	30 0.00%	27 0.00%	24 −4.00%	22 4.76%	20 −4.76%	**370** **3.06%**
2	**1,145** **−0.43%**	**346** **−20.28%**	184 1.66%	102 0.00%	67 −1.47%	49 −2.00%	40 0.00%	37 5.71%	30 0.00%	27 0.00%	25 0.00%	21 0.00%	21 0.00%	**364** **1.39%**
3	1,150 0.00%	435 0.23%	**140** **−22.65%**	102 0.00%	67 −1.47%	49 −2.00%	39 −2.50%	35 0.00%	29 −3.33%	26 −3.70%	24 −4.00%	22 4.76%	20 −4.76%	360 0.28%
4	1,153 0.26%	430 −0.92%	181 0.00%	**76** **−25.49%**	66 −2.94%	52 4.00%	40 0.00%	36 2.86%	30 0.00%	27 0.00%	24 −4.00%	21 0.00%	20 −4.76%	362 0.84%
5	1,151 0.09%	435 0.23%	185 2.21%	102 0.00%	**50** **−26.47%**	50 0.00%	39 −2.50%	35 0.00%	30 0.00%	27 0.00%	25 0.00%	22 4.76%	20 −4.76%	360 0.28%
6	1,153 0.43%	434 0.00%	184 1.66%	102 0.00%	68 0.00%	**37** **−26.00%**	39 −2.50%	34 −2.86%	29 −3.33%	26 −3.70%	24 −4.00%	22 4.76%	20 −4.76%	363 1.11%
7	1,151 0.09%	434 0.00%	183 1.10%	102 0.00%	68 0.00%	50 0.00%	**29** **−27.50%**	35 0.00%	31 3.33%	27 0.00%	24 −4.00%	22 4.76%	20 −4.76%	358 −0.28%
8	1,153 0.26%	433 −0.23%	183 1.10%	103 0.98%	66 −2.94%	50 0.00%	40 0.00%	**25** **−28.57%**	30 0.00%	26 −3.70%	25 0.00%	20 −4.76%	20 −4.76%	360 0.28%
9	1,150 0.00%	430 −0.92%	183 1.10%	102 0.00%	65 −4.41%	50 0.00%	39 −2.50%	35 0.00%	**21** **−30.00%**	27 0.00%	25 0.00%	21 0.00%	19 −9.52%	361 0.56%
10	1,150 0.00%	432 −0.46%	185 2.76%	101 −0.98%	66 −2.94%	51 2.00%	39 −2.50%	36 2.86%	30 0.00%	**19** **−29.63%**	24 −4.00%	21 0.00%	20 −4.76%	358 −0.28%
11	1,151 0.09%	434 0.00%	183 1.10%	104 1.96%	69 1.47%	50 0.00%	40 0.00%	35 0.00%	30 0.00%	27 0.00%	**17** **−32.00%**	21 0.00%	20 −4.76%	360 0.28%
12	1,152 0.17%	436 0.46%	182 0.55%	102 0.00%	68 0.00%	51 2.00%	40 0.00%	35 0.00%	31 3.33%	27 0.00%	24 −4.00%	**14** **−33.33%**	20 −4.76%	360 0.28%
13	1,147 −0.26%	436 0.46%	185 2.76%	102 0.00%	66 −2.94%	50 0.00%	40 0.00%	35 0.00%	31 3.33%	28 3.70%	24 −4.00%	21 0.00%	**14** **−33.33%**	358 −0.28%
Album	**1,210** **5.18%**	**466** **7.45%**	**199** **10.03%**	**112** **9.83%**	**74** **9.18%**	**55** **10.74%**	**45** **10.20%**	**40** **14.17%**	33 9.57%	30 12.59%	27 9.88%	23 8.24%	23 8.24%	**185** **−48.43%**

Notes. In the baseline case, the price of each single is \$0.99, and the price of the album is \$9.99. The first column of the table indicates the price of which song (or the album) is changed. Song price change is from \$0.99 to \$1.29. Album price change is from \$9.99 to \$12.99. The market size is set to 5,000. The first number in each cell is the sales volume; the second number is the percentage change. Bold values are considered significant.

popular songs than they are to the price of less popular songs. In addition, we find that when the price of a song increases, album sales tend to increase, and sales of other songs seem to drop slightly. This is because when the price of any single song increases and the price of the album remains unchanged, the album (bundle) becomes preferable to the individual song(s). As some people move from singles to the album, the sale of other singles within that album should also drop. This is consistent with the results that the sales of all singles increase with an increase in album price.

We also note that when the single song price change occurs to higher-ranked songs, its impact on the sales of other songs and the album is larger. For example, as shown in row 1 in Table 12, the price increase of the most popular song has a positive and significant effect on the sales of album. At the same time, the sales of the second most popular song also drop significantly. However, the price increase of less popular song, say, song 13, will not significantly affect the sales of other singles and the album. The reason is that, from Equations (9) and (10), we know only songs that provide positive utility will affect consumers' choices between single songs and the album. Clearly, songs with higher rankings have higher probabilities to provide positive utilities. Therefore, their price changes will have a larger impact on sales of other songs and the album. On the other hand, when the album price is increased, we see a significant drop in album sales. And the sales of high-ranking songs increase significantly. This is also intuitive because when the album price rises, consumers find it preferable to "cherry pick" the popular singles.

6.5. Optimal Pricing Strategy

Given consumers' average price sensitivity and the distribution of values of songs in a typical album, we can explore a label's "optimal price strategy" under the assumptions above. The probability of purchase generated from the Probit utility function does not have a closed-form solution. Therefore, we rely on numerical methods to solve the optimization problem.

A common assumption made in studies of information goods is that information goods have a marginal cost of zero, and this assumption can be extended to the digital music context. Although Label X spends a large amount of money producing the album, the cost of reproducing digital single songs and albums is close to zero. Given this, maximizing Label X's profit is equivalent to maximizing its revenue.[15] In addition, given that Label X is a monopolist on its songs, the search space for optimal prices should be from 0 to infinity. However, in reality, music labels have to compete against piracy and other consumption channels, which is not captured in our model. To make the simulation condition closer to the reality, we restrict the search space of single songs to the four values: \$0.69, \$0.89, \$1.09, and \$1.29. These four values were selected based on observed prices at various online stores, and thus, although our analysis is representative of existing prices, we could repeat it using any reasonable set of prices in the search space. We keep the relatively wide interval because the computational complexity is increasing rapidly with the number of candidate price points for each song. We also searched various digital music marketplaces to identify the lowest and highest album prices observed and we thus restrict the search space for album prices to \$5 to \$20, with a \$2 interval.

The simulated optimal prices for singles and the album are displayed in Table 13. For a typical album, the price of the first six songs is \$1.29. Songs ranked from 7 to 10 should be priced at \$1.09, whereas the last three songs' optimal price is \$0.89. There is a lot of diversity in the optimal prices of songs with different rankings, and the traditional \$0.99 appears to be below the optimal price for the majority of songs. The optimal album price is \$7, which is lower than the current

Table 13 Optimal Prices (\$)

Rank	Optimized album price[a]	Fixed album price	Rank	Optimized album price	Fixed album price
1	1.29	1.29	9	1.09	1.09
2	1.29	1.29	10	1.09	1.09
3	1.29	1.29	11	0.89	1.09
4	1.29	1.29	12	0.89	0.89
5	1.29	1.29	13	0.89	0.89
6	1.29	1.09	Album	7.00	9.99
7	1.09	1.09	ΔR (%)	18.22	5.79
8	1.09	1.09	ΔCS (%)	23.04	−11.07

[a]We also conducted the same analysis using a finer search grid—song prices were chosen from 0.59, 0.69, 0.79, 0.89, 0.99, 1.09, 1.19, 1.29, and 1.39, and the choices of album price were from \$5 to \$20, with a \$1 interval. We found that the optimal prices identified in this new setting are similar to what we report in this table (difference is within \$0.1), and the optimal album price remains \$7. In that case, ΔR is 19.46%, and ΔCS is 20.07%. This means that the firm can receive only slightly higher revenue by using a finer search grid to identify optimal prices. But the computational burden to find out the slightly better prices for singles and the album is much higher. CS, consumer surplus.

[15] We perform our analysis as if the label and retailer were integrated and analyze the pricing problem from the standpoint of the entire industry. In reality, the retailer keeps some portion of the revenues. When the joint revenue is maximized, both the label's and the retailer's revenue are maximized simultaneously. However, music retailers may incur menu cost when they frequently change music prices. Some music retailers, such as iTunes, are also platform owners and so their objective is not maximizing revenue generated by selling digital music, but maximizing the joint revenue they can generate from selling the platform and digital music. As such, it is possible that the optimal pricing strategy for music retailers is different from that for music labels. Our model could be modified to explore the pricing problem for the retailer alone if we have information on the retailer's cost and revenue structure.

"typical" digital album price of $9.99.[16] This set of prices suggests that a better pricing strategy for an album is to increase the prices of the popular songs and decrease the prices of the least popular songs, while at the same time reducing the price of the album to improve album sales.

The intuition for the increase in popular songs' prices is twofold. First, the demand of high-ranked songs is inelastic, and the additional revenue generated from increased prices can outweigh the loss from reduced sales. Second, higher single song prices will also drive consumers to purchase the album, and thus album sales will increase even at the same album price level. Extending this, our results suggest that digital albums were overpriced at the time of the experiment and thus many consumers preferred singles over the album. The optimal prices are estimated to result in an 18.22% increase in revenue and a 23.04% increase in consumer surplus.

At a high level, excluding menu cost considerations, nonuniform pricing allows labels to procure greater revenues for artists and copyright holders. Surprisingly, under the "optimal" prices, consumer surplus is also higher than current levels, largely because of lower album prices and consumers trading up from purchasing singles to purchasing the album. Since consumer and producer surplus are both higher under the "optimal" prices, social welfare is also higher, with this strategy representing a Pareto-improving scenario. However, labels may also be concerned that when the prices of digital albums decrease, consumers will switch from purchasing physical albums to digital albums. To analyze this possibility, we fix album price at $9.99 to eliminate additional cross-channel cannibalization introduced by a lower digital album price. The optimal prices for singles in this case are similar to those in the first case. Label X can still receive higher revenue (a 5.79% increase), but consumer surplus will decrease. In other words, in the second case, Label X extracts more surplus from the consumers with tiered pricing, but consumers lose surplus relative to uniform $0.99 pricing.

6.6. Album-Only Policy

Currently in the music industry there is significant debate regarding whether labels should sell singles at all. Some practitioners believe that unbundling the album, especially when they know little about how to set the single prices appropriately, may lead to less profit. In addition, some artists insist that an album should be appreciated as a whole. Angus Yong, a guitarist of the Australian heavy metal group AC/DC, was quoted in the London *Telegraph* (2008) as saying "we honestly believe the songs on any of our albums belong together. If we were on iTunes, we know a certain percentage of people would only download two or three songs from the album—and we don't think that represents us musically." Although we cannot measure the loss (if any) in artistic value when the album is unbundled, we can at least show whether Label X can earn higher revenues for its artists and copyright holders compared to the current unbundled policy. Bundling literature has established that if the prices of both bundles and their components are optimized, the mixed bundling strategy is no worse than the pure bundling strategy, because pure bundling is in fact a special case of mixed bundling. However, when the single songs' prices are fixed at $0.99, their prices may not be optimized. Therefore, it is possible that the mixed bundling strategy at wrong prices may do worse than pure bundling at a more appropriate price. To test whether this is the case, we first simulate the revenues Label X receives under different album prices. We try 16 possible price points, from $5 to $20 with $1 intervals. We find that for a typical album, the optimal album price under an album-only policy is $6. In other words, revenue is maximized when the album price is set to $6.[17] We then compare the revenue and consumer surplus under album-only policies (considering both album only at $6 per album and album only at $9.99 per album) with the preexperiment pricing scheme, under which both full-length albums and singles are available and the prices of singles are all $0.99, whereas album prices are $9.99 (Table 14).

Table 14 **Album-Only Policy (Compared to Preexperiment Pricing Strategy)**

Album price	Album price = $9.99 (%)	Album price = $6 (%)
Change in album sales	25.39	151.31
Change in revenue	−20.08	−3.39
Change in consumer surplus	−57.01	2.81

As expected, album sales under the album-only policy are higher than under an unbundled policy, with album sales more than doubling when album price is $6. However, album-only policies result in decreases in revenue. Even if the album price is set optimally at $6, the revenue Label X makes is less than in the preexperiment unbundled strategy, where album

[16] We also note that this optimal digital price is lower than optimal prices for CDs in physical markets. There could be a variety of explanations for this difference. One candidate explanation is that eliminating the manufacturing and distribution costs of physical CDs (estimated by a variety of industry sources at $1–$3) could lower the optimal market price. Another possibility is that the presence of piracy and other consumption channels (e.g., Spotify) have lowered consumers' willingness to pay versus what it would have been in pre-Internet markets. It is also possible that consumers' expectations about what music is "worth" have changed versus pre-Internet markets due to other factors.

[17] In this set of analyses, we do not consider the substitution effect between digital albums and physical albums.

and single song prices were not even set optimally. Consumer surplus is slightly higher under the album-only policy when the album price is $6 (because of the decreased album prices), but much lower when the album price is $9.99.

To conclude, if the album price is set to $9.99, an album-only policy leads to lower revenue and lower consumer surplus compared to the preexperiment pricing strategy ($0.99 for singles and $9.99 for the album). When the album price is set optimally at $6, the revenue and consumer surplus are just comparable to those in the preexperiment pricing strategy. Given that the optimal tiered pricing strategy outperforms the preexperiment uniform pricing strategy in terms of revenue and consumer surplus (as discussed above), an album-only strategy must be less preferred than the optimal tiered pricing strategy. This is also consistent with bundling theory that a mixed bundling strategy (under optimized prices) is at least as profitable as is a pure bundling strategy. Our simulation results suggest that unbundling albums does lead to lower album sales. However, the joint revenue from both singles and albums will increase. If the artistic decision is to sell the album as an integral work, this will result in lower overall revenues. If the objective is to maximize revenues to the artist and copyright holders, unbundling achieves a better result. Unbundling also results in engagement with a larger number of consumers because some consumers will only buy individual songs, but not the whole album.

6.7. Robustness Check

In the main model, we assume that ε_{iajt} is i.i.d. across songs. Here, we examine how different distributional assumptions on the correlation structure affect the estimation results.

Identifying the correlation across songs within an album is not intuitively obvious. To see this, note that a positive correlation between songs suggests that if users prefer one song, they are also more likely to prefer other songs. In other words, all else equal, albums become more preferable than singles when correlation is large. Thus, an increase in the price of a song would lead to a larger shift toward the album. The following robustness checks confirm our intuition.

For simplicity, we assume that the covariance between the idiosyncratic preference shocks for every pair of singles is the same (i.e., all off-diagonal elements in the covariance matrix are the same) and less than 1 (as we assume the variance of these shocks equals 1). Estimating the full variance–covariance matrix is infeasible, because the number of parameters needed to estimate will be $n*(n-1)/2$, where n represents the number of songs in the album, making this intractable. In addition, different albums in our sample have different n's, and thus the dimension of the variance-covariance matrix is different for different albums. Finally, as we discussed in §5, this variance–covariance structure is equivalent to assuming heterogeneous consumer tastes. Even with this simplifying assumption, as we noted, we will not be able to estimate the correlation parameter, because we allow for a free album value parameter, which cannot be jointly identified with the correlation parameter. The best thing we can do is to try different correlation parameters and see which parameters give us the results that match the data patterns best. Including the main model, we tried three different levels of correlation: 0, 0.5, and 0.9. In theory, we could try more levels of correlation. However, it is extremely time consuming to do so. We believe that the three correlation level we test here can give us a good idea how the correlation assumption affects the estimation results.

Table 15 **Reduced-Form and Structural Price Elasticity Estimates (Different Correlations)**

Rank	Reduced form (%)	Structural (correlation = 0) (%)	Structural (correlation = 0.5) (%)	Structural (correlation = 0.9) (%)
1	−13.0	−15.3	−15.4	−16.44
2	−7.2	−20.3	−21.14	−23.02
3	−16.0	−22.7	−24.19	−26.34
4	−8.2	−25.5	−26.61	−29.13
5	−13.3	−26.5	−27.54	−31.88
6	−20.6	−26.0	−27.45	−30.91
7	−17.6	−27.5	−30.95	−34.09
8	−26.4	−28.6	−32.35	−34.21
9	−21.7	−30.0	−32.26	−33.33
10	−27.0	−29.6	−34.48	−35.48
11	−16.1	−32.0	−34.62	−34.48
12	−15.0	−33.0	−34.33	−36.36
13	−25.1	−33.0	−34.78	−40.94
Album	−53.55	−48.43	−36.23	−30.36

Table 15 summarizes the comparison between the reduced-form and structural (under different correlation structures) estimates of the price elasticities. The results suggest the following: (1) As correlation increases, singles' own-price elasticity increases faster with rank. (2) As the correlation increases, own-price elasticity of albums is smaller. (3) The elasticity estimates we get under low correlation assumptions are more comparable to the range of the reduced-form estimates. Although the reduced-form estimates are not a perfect benchmark to be compared with, because the effects of concurrent price changes are often entangled, the overall range of the reduced-form estimates of price elasticity should still provide a good reference point, given the nature of the price experiment. The structural estimates under high correlation assumptions are too different from the reduced-form results, and thus inconsistent with the data observation. Therefore, we conclude that low correlation assumption is more reasonable.

6.8. Stratified Analysis

So far, we have used all of our data to estimate consumers' price sensitivity and the value distribution of the singles in a typical album. One concern with this approach is that consumers' price sensitivity and the value distribution of singles in an album may differ across genre, popularity, and some other album characteristics. In this section, we briefly discuss how the own- and cross-price elasticity and optimal prices change across different subsamples.

To do this we first select the three genres with largest number of albums in our data set: pop/rock, alternative, and country. We find that the estimated price coefficient is higher in terms of absolute value for alternative and country versus pop/rock. We also see that consumers value alternative albums the highest, followed by pop/rock and then country. Therefore, a higher percentage of potential consumers of alternative music would purchase the whole album, whereas potential consumers of pop/rock tend to purchase hit songs more frequently. In addition, we find that demand elasticity is lower for pop/rock songs, and thus the optimal single prices for pop/rock are generally higher than the optimal single prices for the other two genres. In all three cases, consumer surplus increases when the tiered pricing strategy is implemented and album prices are reduced to optimal levels.

We then stratify our sample by popularity. To be more specific, we rank the all the albums by the maximum single sales that occurred prior to the experiment, and define the top one-third albums as "popular" albums, the middle one-third albums as "average" albums, and the bottom one-third albums as "unpopular" albums. We find that consumers' price sensitivity and song and album values for popular and average albums are quite similar. Consumers are more sensitive to price changes in unpopular albums, and the single song values and album values are lower for unpopular albums. In addition, since the album value of long tail music is low relative to single song values, when the prices of single songs increase, most consumers choose not to buy anything, and only a few switch to purchasing the album. The optimal pricing strategies for the popular and average songs are similar, and the optimal single prices and album prices for unpopular songs are all lower. Likewise, revenue increases under a tiered pricing strategy with lower album price. Again, under the tiered pricing strategy and reduced album prices, consumer surplus is higher, with the increase in consumer surplus being most significant in the "unpopular" category.

7. Discussion

The combined effects of greater choice of wholesale price points and the unbundling of singles from albums in digital markets has made understanding consumer demand a significant priority for music labels. However, there have been few opportunities for the music industry to study the effects of price changes because of the highly structured pricing models of the major online digital music stores and because the pricing problem is complex, requiring advanced analytical tools that account for the relationship between songs and their parent albums.

Digitalization of music and the ability to use "agency pricing" to directly set retail prices for their products have given labels to the opportunity to study consumer demand characteristics and guide their pricing strategies. Our study is the first we are aware of to leverage the large data sets that are now available to firms in the context of advanced data analytics, structural modeling, and price experimentation available in digital markets. The result is that we can model the complexities of consumer demand for songs and albums to provide optimized pricing guidance while also exploring producer and consumer surpluses under counterfactual pricing policies. Our methodology could be used as a template to analyze pricing questions for other firms selling digital goods, particularly when those goods may be bundled.

Our estimation results show a large heterogeneity in values across songs in an album and heterogeneity that decreases with song popularity within the album. The own-price elasticity of demand is higher for less popular songs. We also find that changes in the prices of singles or albums create complicated substitution effects between singles and their parent albums. When single songs' prices, especially those of popular singles, increase, their own sales will drop, album sales will increase, and sales of some other singles in the same album may drop. All of this is consistent with existing theory on bundling.

Together we find that tiered pricing strategies can significantly improve Label X's joint profit from single song sales and album sales. We find that consumer surplus also increases when a tiered pricing strategy is implemented together with optimal reduced album prices. We show that, in general, higher-ranked songs should be priced higher, whereas lower-ranked songs should be priced lower. We also find that albums' own-price elasticity is high, and thus Label X should reduce album prices (which averaged $9.99 at the time of the experiment). Finally, we show that increasing prices of single songs can encourage consumers to purchase the album.

Our "album-only" policy simulation also shows that even under optimal album pricing, a pure bundled "album-only" policy cannot bring higher revenue than mixed bundling with uniform pricing, and brings less revenue and consumer surplus than mixed bundling with optimized tiered pricing. This suggests that, in the absence of an artistic priority to maintain albums as integral works, unbundling digital albums into mixed bundles (singles and albums) is both an effective policy to increase revenue from sales to producers and can represent a Pareto-improving outcome to consumers.

As with all empirical work, our study has several limitations. First, because of data limitations, we have not included consumer heterogeneity in terms of their price sensitivity. Since our aggregate data only include changes from one particular price point to another ($0.99 to $1.29), we cannot separately identify parameters associated with heterogeneity. Second, because of tractability issues, we cannot estimate the true correlation structure among individuals' preference shocks to different single songs. Instead, we can only use different simplified correlation structures to estimate the model and see which assumptions generate results that are closest to the observed data patterns. Third, because we only have two price points in our data in our elasticity analyses and pricing experiments, we must assume that reactions to the observed price increase from $0.99 to $1.29 are representative of other possible price regions. Fourth, our analysis does not consider the impact of these digital price changes on demand for physical CDs.[18] Finally, our estimates are driven by price changes made according to a rule rather than a purely random trial, and although we have given evidence that we are able to exploit exogenous price variation in the data, we acknowledge that the gold standard for price experimentation would be a randomized controlled trial.

In spite of these limitations, we believe that our results reveal important insights about the managerial importance and advantages of using advanced data analytics techniques to analyze optimal pricing and bundling strategies for creative works sold in digital marketplaces. In addition to advancing the managerial understanding of these important issues, our work advances the academic literature by being one of the first papers to apply a structural framework to real-world data to analyze bundling of information goods. Our work is also one of the first examples of an academic–industry partnership to develop and apply "big data" management paradigms to a pricing problem in the media industry—and our results guided digital pricing decisions at the major label after the experiment was completed. We hope that our work can pave the way for the future research in this important area as well as future partnerships between academics and practitioners interested in applying a "big data" approach to management.

[18] However, we point out that the extant literature strongly suggests that marketing changes in digital channels do not affect short-term demand in physical channels. For example, Biyalogorsky and Naik (2003) find that Tower Records' addition of an Internet distribution channel did not significantly cannibalize its retail sales, Waldfogel (2009) finds that authorized YouTube viewing of television content has only a very small net displacement effect on over-the-air viewing, Danaher et al. (2010) find no change in demand for NBC's DVD content at Amazon.com associated with NBC's closing or reopening of its digital distribution channel on iTunes, and Hu and Smith (2013) find that delaying the availability of Kindle eBooks has no statistical impact on demand for hardcover books.

Acknowledgments

All authors contributed equally and are presented in alphabetical order. The authors thank an anonymous record label for providing the data necessary to conduct this research. The authors acknowledge Carnegie Mellon's iLab for generous financial support. The authors thank the Center for the Analysis of Property Rights and Innovation at University of Texas at Dallas for research support. Brett Danaher acknowledges support from the National Bureau of Economic Research through the Economics of Digitization and Copyright grant program.

References

Adams WJ, Yellen JL (1976) Commodity bundling and the burden of monopoly. *Quart. J. Econom.* 90(3):475–498.

Armstrong M (1996) Multiproduct nonlinear pricing. *Econometrica* 64(1):51–75.

Bajari P, Hortaçsu A (2003) The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND J. Econom.* 34(2):329–355.

Bakos Y, Brynjolfsson E (1999) Bundling information goods: Pricing, profits, and efficiency. *Management Sci.* 45(12):1613–1630.

Belinfante A, Johnson RL (1982) Competition, pricing and concentration in the U.S. recorded music industry. *J. Cultural Econom.* 6(2):11–24.

Berry S (1994) Estimating discrete-choice models of product differentiation. *RAND J. Econom.* 25(2):242–262.

Berry S, Levinsohn J, Pakes A (1995) Automobile prices in market equilibrium. *Econometrica* 63(4):841–890.

Berry S, Levinsohn J, Pakes A (2004) Differentiated products demand systems from a combination of micro and macro data: The new car market. *J. Political Econom.* 112(1):68–105.

Biyalogorsky E, Naik P (2003) Clicks and mortar: The effect of online activities on off-line sales. *Marketing Lett.* 14(1):21–32.

Breidert C, Hahsler M (2007) Adaptive conjoint analysis for pricing music downloads. *Adv. Data Anal.* 4:409–416.

Brynjolfsson E, Dick AA, Smith MD (2010) A nearly perfect market? Differentiation versus price in consumer choice. *Quant. Marketing Econom.* 8(1):1–33.

Burstein ML (1960) The economics of tie-in sales. *Rev. Econom. Statist.* 420(1):68–73.

Chintagunta PK (2001) Endogeneity and heterogeneity in a probit demand model: Estimation using aggregate data. *Marketing Sci.* 20(4):442–456.

Chu CS, Leslie P, Sorensen A (2006) Good enough bundling. Working paper, Stanford University, Stanford, CA.

Chuang J, Sirbu M (1999) Optimal bundling strategy for digital information goods: Network delivery of articles and subscriptions. *Inform. Econom. Policy* 11(2):147–176.

Crawford GS (2008) The discriminatory incentives to bundle: The case of cable television. *Quant. Marketing Econom.* 6(1):41–78.

Crawford GS, Yurukoglu A (2012) The welfare effects of bundling in multichannel television markets. *Amer. Econom. Rev.* 102(2): 643–685.

Danaher B, Dhanasobhon S, Smith MD, Telang R (2010) Converting pirates without cannibalizing purchasers: The impact of digital distribution on physical sales and internet piracy. *Marketing Sci.* 29(6):1138–1151.

Davis P (2006) Spatial competition in retail markets: Movie theaters. *RAND J. Econom.* 37(4):964–982.

Derdenger T, Kumar V (2013) The dynamic effects of bundling as a product strategy. *Marketing Sci.* 32(6):827–859.

Elberse A (2010) Bye-bye bundles: The unbundling of music in digital channels. *J. Marketing* 74(3):107–123.

Geng X, Stinchcombe M, Whinston A (2005) Bundling information goods of decreasing value. *Management Sci.* 51(4):662–667.

Hann I, Terwiesch C (2003) Measuring the frictional cost of online transactions: The case of a name-your-own-price channel. *Management Sci.* 49(11):1563–1579.

Hitt L, Chen P (2005) Bundling with customer self-selection: A simple approach to bundling low-marginal-cost goods. *Management Sci.* 51(10):1481–1493.

Hu YJ, Smith MD (2013) The impact of ebook distribution on print sales: Analysis of a natural experiment. Working paper, Carnegie Mellon University, Pittsburgh, PA.

McAfee PP, McMillan J, Whiston MD (1989) Multiproduct monopoly, commodity bundling, and correlation of values. *Quart. J. Econom.* 104(2):371–383.

Nevo A (2000) Mergers with differentiated products: The case of the ready-to-eat cereal industry. *RAND J. Econom.* 31(3):395–421.

Nevo A (2001) Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69(2):307–342.

Shiller B, Waldfogel J (2009) Music for a song: An empirical look at uniform song pricing and its alternatives. Working paper, University of Minnesota, Minneapolis.

Smith E (2009) Apple changes tune on music pricing. *Wall Street Journal* (January 7), http://online.wsj.com/article/SB123126062001057765.html.

Stigler GJ (1963) United States v. Loew's Inc.: A note on block booking. *Supreme Court Rev.* 152–157.

Telegraph (London) (2008) iTunes under threat as bands take their business elsewhere. (September 25), http://www.telegraph.co.uk/culture/music/3561194/iTunes-under-threat-as-bands-take-their-business-elsewhere.html.

Waldfogel J (2009) Lost on the web: Does web distribution stimulate or depress television viewing? *Inform. Econom. Policy* 21(2):158–168.

Zhu K, MacQuarrie B (2003) The economics of digital bundling: The impact of digitization and bundling on the music industry. *Comm. ACM* 46(9):264–270.